

Keith Steward, Ph.D.

Machine Learning Engineering Leader | Production ML Systems Builder

Boston, MA | 508-318-8035 | keith.steward@gmail.com

[LinkedIn](#) | IntelligentDataWorks.ai | [IDW LinkedIn](#)

PROFILE

Technical AI/ML engineering leader who builds enterprise GenAI platforms and the teams that operate them in regulated environments. Combines hands-on Python, APIs, SQL, RAG, vector databases, agentic orchestration, Kubernetes, MLOps, and CI/CD with 20 years of people leadership. Built an 11-engineer organization that delivered 70+ production AI automations and previously advised 200+ AWS enterprises. Brings the architecture, governance, executive partnership, and delivery discipline to turn finance processes into secure, measurable AI capabilities.

SELECTED QUALIFICATIONS

- **AI Engineering Leadership:** team building, mentoring, Agile execution, architecture direction, engineering standards, and delivery accountability
- **Enterprise GenAI:** LLM applications, RAG, vector databases, AI agents, LangChain, LangGraph, evaluation, observability, and workflow orchestration
- **Cloud & Full-Stack Delivery:** Python, FastAPI, REST APIs, SQL, PostgreSQL, AWS Bedrock, SageMaker, Kubernetes, Docker, and distributed systems
- **Governance & Operations:** responsible AI, privacy, security, auditability, data governance, MLOps, DevOps, CI/CD, model transparency, bias, and drift controls
- **Executive Enablement:** finance and regulated-industry roadmaps, use-case prioritization, ROI framing, stakeholder advisory, and AI Center of Excellence leadership

PROFESSIONAL EXPERIENCE

Technical Co-Founder & CEO | Intelligent DataWorks | Boston, MA | Aug 2023 to Present

Lead architecture, engineering, and product execution for an applied AI SaaS company while remaining hands-on in platform development.

- Architected and helped build AssistX, a production GenAI platform combining Python, FastAPI, PostgreSQL/PGVector, LLMs, AI agents, RAG, LangChain, LangGraph, APIs, and AWS services.
- Recruited, trained, and technically led an 11-engineer distributed organization across AI/ML, full-stack software, cloud, and DevOps through Agile sprints, code reviews, testing, releases, and operations.
- Delivered the initial production platform within eight months, supporting 70+ AI automations across more than 20 workflows and accelerating selected work by up to 50x.
- Established reusable controls for prompt safety, input validation, human review, evaluation, observability, data isolation, security, and multi-tenant SaaS operation.
- Initiated strategic acquisition discussions regarding IDW's AI platform as part of a planned transition to a senior AI/engineering leadership role.

Senior Director of Engineering | SuccessKPI | Remote | Jun 2022 to Aug 2023

Directed engineering and architecture for the Workforce Management component of an AI-powered enterprise contact-center SaaS platform.

- Recruited and led a seven-person distributed engineering team delivering forecasting, optimization, scheduling, analytics, APIs, and automated communications on AWS.
- Translated detailed requirements into architecture, algorithms, data flows, Python services, work breakdown, testing, CI/CD, and iterative production releases.
- Delivered the first production release within six months and onboarded enterprise customers, including the U.S. Department of State, while maintaining quality and operational reliability.

Senior Director, Data Analytics & Operations | TriMark USA | Mansfield, MA | Aug 2021 to May 2022

Owned enterprise data modernization for a \$1B+ organization integrating 11 acquired businesses and heterogeneous ERP systems.

- Partnered with Finance, Marketing, Supply Chain, and technology executives to gather and prioritize 50+ analytics use cases and KPIs, define an AWS roadmap, and secure executive support and funding.

- Upskilled a four-person analytics group and grew it into a seven-member AWS engineering and analytics function with planning, hiring, mentoring, budget, risk, and delivery accountability.
- Delivered the company's first AWS data lake within seven months and reduced cross-ERP reporting from roughly two months to one day, a 60x improvement in insight delivery.

Technical Lead, Cloud Data Privacy & AI Governance | MassMutual | Boston, MA | Nov 2020 to Aug 2021

Established cloud-native engineering across five Data Privacy, AI Governance, and Responsible AI initiatives in regulated financial services.

- Led architecture and hands-on implementation for Privacy Rights Management, enterprise metadata, governance APIs, and policy automation using secure AWS patterns.
- Partnered with Enterprise Architecture and Site Reliability Engineering to define security, IAM, Kubernetes, serverless, infrastructure-as-code, testing, and CI/CD standards.
- Advised AI Governance & Ethics on MLOps, model transparency, bias, drift, monitoring, auditability, and SageMaker-based lifecycle controls.
- Delivered the first cloud-native privacy-rights platform within eight months, automating approximately 90% of a formerly manual workflow supporting CCPA, CPRA, and GDPR obligations.

VP of AI/ML | NorthBay Solutions | Boston, MA | Jan 2019 to Nov 2020 | (Role ended in company-wide COVID layoff).

Recruited out of AWS by the NorthBay Solutions CEO to establish a new enterprise AI/ML practice and reusable production data-and-ML capability.

- Proposed an AWS Data Lake and ML accelerator, built the initial platform hands-on using Python and AWS, and led a distributed engineering team toward customer demonstrations and delivery.
- Advised enterprise clients in financial services, higher education, and data services on AI strategy, solution architecture, data lakes, knowledge graphs, and production machine learning.

Senior Specialist Solutions Architect, AI/ML & Big Data | Amazon Web Services | Boston, MA | May 2015 to Dec 2018

Advised enterprises and AWS technical teams on production machine learning, distributed analytics, data platforms, and cloud adoption.

- Advised 200+ enterprise customers across financial services, healthcare, telecommunications, manufacturing, and technology on AI roadmaps, SageMaker, TensorFlow, distributed data, security, and production architecture.
- Founded and led AWS technical communities for Big Data and AI/ML, growing the AI/ML community to 80+ global specialists through standards, mentoring, workshops, and reusable technical assets.
- Created 170+ enablement assets, trained 180+ solution architects, and delivered 50+ technical presentations and workshops on AI/ML and distributed data systems.

EARLIER RELEVANT EXPERIENCE

- Healthcare and regulated-software experience includes leading a five-engineer team, including a sub-team manager, that delivered the HIPAA-governed, FDA-cleared Lythos dental scanning platform and its fault-tolerant AWS cloud service.
- Earlier engineering leadership includes direct management of a 25-person distributed SaaS team at CambridgeSoft and scientific software development spanning commercial life-sciences platforms.

EDUCATION & RECOGNITION

Ph.D., Molecular Biology / Microbiology & Immunology, University of Western Ontario. AWS ML Elite Gold and Platinum; AWS re:Invent Rockstar; named inventor on multiple patents.

TECHNOLOGIES

LLMs • RAG • AI agents • agentic workflows • Python • REST APIs • SQL • PostgreSQL • PGVector • vector databases • Kubernetes • Docker • AWS Bedrock • SageMaker • Vertex AI • LangChain • LangGraph • knowledge graphs • AWS Neptune • MLOps • CI/CD • DevOps • data governance • responsible AI • model evaluation • model transparency • bias detection • drift detection • observability • FastAPI • Spark • PySpark • data lakes • data pipelines • distributed systems • microservices • EventBridge • API Gateway • ECS/Fargate • AWS CDK • Git • Linux • TensorFlow • machine learning • predictive analytics